Assessing the Value of Frost Forecasts to Orchardists: A Dynamic Decision-Making Approach[1]

RICHARD W. KATZ AND ALLAN H. MURPHY

Department of Atmospheric Sciences, Oregon State University, Corvallis 97331

ROBERT L. WINKLER

Graduate School of Business, Indiana University, Bloomington 47405

(Manuscript received 11 May 1981, in final form 26 January 1982)

ABSTRACT

The methodology of decision analysis is used to investigate the economic value of frost (i.e., minimum temperature) forecasts to orchardists. First, the fruit-frost situation and previous studies of the value of minimum temperature forecasts in this context are described. Then, after a brief overview of decision analysis, a decision-making model for the fruit-frost problem is presented. The model involves identifying the relevant actions and events (or outcomes), specifying the effect of taking protective action, and describing the relationships among temperature, bud loss, and yield loss. A bivariate normal distribution is used to model the relationship between forecast and observed temperatures, thereby characterizing the quality of different types of information. Since the orchardist wants to minimize expenses (or maximize payoffs) over the entire frost-protection season and since current actions and outcomes at any point in the season are related to both previous and future actions and outcomes, the decision-making problem is inherently dynamic in nature. As a result, a class of dynamic models known as Markov decision processes is considered. A computational technique called dynamic programming is used in conjunction with these models to determine the optimal actions and to estimate the value of meteorological information.

Some results concerning the value of frost forecasts to orchardists in the Yakima Valley of central Washington are presented for the cases of red delicious apples, bartlett pears, and elberta peaches. Estimates of the parameter values in the Markov decision process are obtained from relevant physical and economic data. Twenty years of National Weather Service forecast and observed temperatures for the Yakima key station are used to estimate the quality of different types of information, including perfect forecasts, current forecasts, and climatological information. The orchardist's optimal actions over the frost-protection season and the expected expenses associated with the use of such information are determined using a dynamic programming algorithm. The value of meteorological information is defined as the difference between the expected expense for the information of interest and the expected expense for climatological information. Over the entire frost-protection season, the value estimates (in 1977 dollars) for current forecasts were \$808 per acre for red delicious apples, \$492 per acre for bartlett pears, and \$270 per acre for elberta peaches. These amounts account for 66, 63, and 47%, respectively, of the economic value associated with decisions based on perfect forecasts. Varying the quality of the minimum temperature forecasts reveals that the relationship between the accuracy and value of such forecasts is nonlinear and that improvements in current forecasts would not be as significant in terms of economic value as were comparable improvements in the past.

Several possible extensions of this study of the value of frost forecasts to orchardists are briefly described. Finally, the application of the dynamic model formulated in this paper to other decision-making problems involving the use of meteorological information is mentioned.

1. Introduction

Meteorological information may be of value in a variety of different situations, and meteorologists and others are increasingly recognizing the need for quantitative assessments of the value of such information. In order to obtain reliable and credible estimates of the value of meteorological information, it generally is desirable to model the decision-making and information-processing procedures adopted by individual users of the information. Decision analysis (Raiffa, 1968; Lindley, 1971; Anderson *et al.*, 1977; Keeney, 1982) represents a set of concepts and procedures for modeling and analyzing complex decision-making problems. Moreover, a methodology for determining the *ex ante* value of information plays a natural and central role in decision analysis. Thus, decision analysis appears to offer a particularly attractive framework within which to study the value of meteorological information.

[1] Supported in part by the National Science Foundation (Atmospheric Sciences Division) under grant ATM80-04680.

0021-8952/82/040518-14\$07.50

A decision-making problem in which meteorological data may be of significant economic value is the so-called "fruit-frost situation." In this situation, fruit growers use weather forecasts and other information to decide whether or not to protect their orchards against freezing temperatures. In recognition of the important role that meteorological information plays in this situation, the National Weather Service (NWS) has developed a fruit-frost program in which orchardists in many regions of the United States are routinely provided with specialized minimum temperature forecasts during those periods or seasons when frost presents a significant hazard. The purpose of this paper is to describe some results of a study in which the methodology of decision analysis was used to investigate the economic value of frost forecasts to orchardists in the Yakima Valley of central Washington.

It is assumed that the goal of an orchardist is to minimize the total expected expense (or, equivalently, maximize the total expected return) over the entire frost-protection season. Because an orchardist's decision as to whether or not to take protective measures on a particular occasion depends on similar decisions and frost events on previous and subsequent occasions, this decision-making problem is *dynamic* in nature. A particular class of dynamic models known as Markov decision processes (e.g., Howard, 1960; Ross, 1970) provides a suitable representation of the fruit-frost situation. Such a dynamic decision-making approach yields more reliable estimates of the value of minimum temperature forecasts to orchardists than an approach based on a static model.

In this paper, we summarize the results of a study in which Markov decision processes were used to estimate the economic value of different types of frost forecasts to fruit growers in the Yakima Valley. The fruit-frost situation and previous studies of the value of minimum temperature forecasts in this situation are discussed briefly in Section 2. Section 3 contains an overview of decision analysis and a description of the dynamic model used to study the fruit-frost problem. The discussion of the model includes treatments of: 1) the relevant actions, events, and payoffs; 2) the models and types of information or forecasts; and 3) Markov decision processes and value-of-information expressions and relationships. Some results of the application of this model to the fruit-frost problem are presented in Section 4, with particular emphasis on the orchardist's optimal actions and the value of different types of forecasts. Section 5 consists of a brief summary and some concluding remarks.

2. The fruit-frost situation and previous studies of the value of frost forecasts

a. *The fruit-frost decision-making situation*

Deciduous fruit trees become particularly susceptible to damage from freezing temperatures, or frost, in the spring when the buds begin to develop into blossoms. Frost can damage or kill buds, resulting in the loss of fruit yield, and extremely low temperatures can even damage the trees themselves. The extent of bud loss for a given temperature is highly dependent on the type and variety of fruit and on the stage of development of the buds. The relationship between bud loss and fruit loss is complex, because considerable bud loss can occur without any corresponding loss in fruit yield. In the Yakima Valley of central Washington, the frost-protection season extends from mid-March through May.

To minimize the damage caused by freezing temperatures, many fruit growers employ devices such as heaters, wind machines, and overhead sprinklers to protect their orchards (e.g., Bagdonas *et al.*, 1978). Prior to 1970, heaters designed to raise the air temperature in orchards were the principal means of protection against frost. However, substantial rises in fuel and labor costs have greatly increased the expenses associated with the use of orchard heaters in recent years. Moreover, the introduction of local air quality standards has placed some restraints on the use of heaters under certain meteorological conditions. As a result, alternative methods of protection have been adopted by many fruit growers in place of (or in addition to) heaters, most notably wind machines and overhead sprinklers. Wind machines generally improve the circulation of air in the orchard and inhibit the formation of a cold layer of air near the surface. Overhead sprinklers add a layer of water to exposed surfaces on the trees, and the latent heat of fusion released when this water freezes maintains the buds (blossoms/fruit) at or slightly above the freezing level. Some risks are involved in the use of these alternatives, however, since colder air may be drawn down from aloft by wind machines under certain conditions and ice accumulations from water freezing over an extended period of time occasionally can cause extensive damage to fruit trees. Some orchardists employ protective devices in combination—for example, heaters and wind machines—thereby increasing the protective effect and/or reducing the likelihood of any adverse impacts. A substantial fraction of the total acreage of fruit trees in the Yakima Valley can be protected from freezing temperatures by employing one or more of these types of protective devices.

Each evening during the frost-protection season a fruit grower in the Yakima Valley must decide whether or not to protect his orchard in the face of uncertainty concerning the actual minimum temperature for that night. Since the use of protective devices involves considerable expense, it may not be reasonable from an economic standpoint to use these devices whenever the slightest chance of frost exists. To aid the grower in making his decision, NWS forecasters at the Yakima Weather Service Office make minimum temperature forecasts for certain key locations in the area and disseminate these forecasts to the orchardists each evening. An orchardist's decision as to whether or not to protect is based at least to some extent on these forecasts, since to be effective the protective measures generally must be taken before the critical temperatures for bud loss are actually reached. Roughly speaking, the value of the minimum temperature forecasts to a fruit grower is measured by the ability of the forecasts to reduce the uncertainty under which the decision is made, resulting in turn in a reduction in the orchardist's expected expenses.

b. *Previous fruit-frost value-of-information studies*

Two previous studies of the value of minimum temperature forecasts to orchardists in the fruit-frost situation have been undertaken within a decision-analytic framework. Baquet *et al.* (1976) investigated the value of frost forecasts to pear growers in the vicinity of Medford, Oregon using an approach based on Bayesian decision making under uncertainty. As the authors indicated, this approach yields an *ex ante* assessment of the value of the relevant information; that is, the expected value of the information before the forecasts and observations actually are obtained (for a brief discussion of the *ex ante* and *ex post* approaches to the assessment of the value of information, see Section 3a). Baquet *et al.* formulated a decision-making model based on the assumption that the orchardist wants to maximize expected return or payoff on a particular occasion, taking due account of his current situation (which, in turn, depends on previous decisions and frost events). Their solution to the value-of-information problem involved the use of a simulation procedure that employed historical NWS data concerning forecasts and observations at Medford, Oregon to generate typical frost seasons. Relative to the value of forecasts based solely on climatological probabilities, Baquet *et al.* estimated that the value of the specialized NWS minimum temperature forecasts was $4.73 per acre per day (averaged over the eight orchardists considered in the study), or approximately $285 per acre over the two-month frost-protection season.

More recently, Katz and Murphy (1979) have studied the value of frost forecasts to orchardists in the Yakima Valley of central Washington. In their initial study, the authors employed a static decision-making model in which the relationships between decisions and frost events on different occasions are ignored. Under the assumption that the goal of the orchardist is to minimize the expected expense on a particular occasion, and ignoring any previous bud loss that has been incurred, such a model can be used to estimate the value of forecasts at a particular time during the frost season. In this regard, Katz and Murphy estimated that the *ex ante* value of the NWS minimum temperature forecasts for the Yakima key station averaged $24 per acre per day over the season (relative to the value of forecasts based on climatological probabilities). This value estimate varied from less than $10 per acre per day at the beginning and end of the season to over $40 per acre per day in the middle of the ten-week frost-protection season. Katz and Murphy emphasized that this static model could not be used to determine the value of these forecasts when the goal of the orchardist is to minimize the expected expense over the entire season.

The model to be described in this paper can be viewed as an extension or generalization of the models employed by Baquet *et al.* (1976) and Katz and Murphy (1979). In formulating this model we have assumed that the orchardist wants to minimize the expected expense (or maximize the expected payoff) over the entire frost-protection season. This objective leads to the development of a dynamic model in which decisions on any particular occasion are related to both past and future decisions and frost events. This model is a member of a class of dynamic models known as Markov decision processes, and an exact numerical solution to this decision-making problem is obtained via dynamic programming. Moreover, a model of temperature forecasts and observations is formulated, and this model provides a basis for determining the value of current NWS minimum temperature forecasts as well as the value of forecasts of improved quality. The dynamic decision-making model and the model of meteorological information are described in greater detail in Section 3.

3. A dynamic decision-making model

In this section we first present a brief overview of decision analysis, which provides the general framework within which this study is conducted. Second, we describe the specific objectives, actions, events, and consequences that constitute the fruit-frost decision-making problem addressed in this paper. Third, different types of meteorological information are considered, ranging from climatological information to perfect forecasts, and models for determining probability distributions that characterize the uncertainty associated with this information are presented. Fourth, and finally, the decision-making problem that confronts the orchardist is formulated as a Markov decision process, and expressions are given for the value of several different types of meteorological information.

a. *A brief overview of decision analysis*

Decision analysis can be defined as a set of concepts and procedures for analyzing decision-making problems involving uncertainty in a logical, rational manner (e.g., Keeney, 1982). For the purposes of this discussion, decision analysis can be considered to consist of the following steps: a) structuring the decision-making problem; b) determining the utilities of the consequences to the decision maker; c) assessing the probabilities of the events; d) evaluating and comparing the alternative actions; and e) estimating the value of information (concerning the events) to the decision maker. Structuring the problem involves identifying the possible actions, events, and consequences (each action–event pair leads to a distinct consequence). If the problem involves a sequence of actions and events, they must be ordered appropriately in the decision-making model.

The utilities express the decision maker's preferences for the consequences. When all the consequences are expressed in monetary terms, it is necessary to assess a utility function for money. Frequently, it is possible to use a model (e.g., a linear model, an exponential model, or a logarithmic model) to represent a decision maker's utility function. A linear model, which implies risk neutrality in the relevant range of monetary payoffs, provides a reasonable first approximation to the decision maker's utility function in certain situations. The use of a linear model also simplifies matters considerably by making it possible to work with expected monetary payoffs instead of expected utilities.

In the third step, the probabilities of the events are determined by combining information from objective sources (e.g., statistical models, historical data) and information in the form of subjective judgments. This process may involve the use of probabilistic models to describe the joint and/or conditional distributions of various quantities. Bayes' theorem provides the framework within which information in the form of likelihoods provided by these models can be combined with prior probabilities to produce posterior probabilities of the relevant events. In a meteorological context, the prior probabilities might be based on climatological data and the likelihoods might be obtained from forecasting procedures or from the results of previous forecasting experiments or programs.

In decision analysis, the alternative actions are evaluated by comparing their expected utilities. The expected utility of an action is the weighted average of the utilities of the consequences associated with that action, where the weights are the probabilities of the relevant events. The preferred action is the action with the highest expected utility. In the case of a linear utility function, choosing the action with the highest expected monetary payoff is equivalent to choosing the action with the highest expected utility.

The value of information plays an important and central role in decision analysis. In this context, the value of a certain type of information in a particular decision-making problem is defined as the maximum amount that the decision maker should be willing to pay for this information. This amount depends on the nature of the problem, on the decision maker's initial state of uncertainty, and on the expected quality of the information. Moreover, information has value only if it can lead to a change in the decision. Thus, the value of information is the cost of information that makes the prior expected utility and the "average" posterior expected utility exactly equal (or, in the case of a linear utility function, the difference in expected monetary payoff before and after the receipt of the information). It is important to note that this analysis is conducted in an *ex ante* or expected value sense; that is, it is an analysis of the expected value of information before that information actually is available. Finally, although perfect information is seldom available, the value of perfect information provides a useful upper bound on the value of all types of imperfect information.

Decision analysis is prescriptive rather than descriptive in nature. That is, the aim of decision analysis is not to describe how decisions are actually made and how information is used in practice, but to prescribe how decisions can be made and information can be used in a rational, optimal manner. Of course, it may be of considerable interest to investigate how the decision maker actually makes decisions and uses information and to compare these procedures with the recommendations forthcoming from a formal prescriptive analysis of the decision-making problem.

b. *Objectives, actions, events, and consequences*

The orchardist is assumed to want to minimize the total expected expense—or, equivalently, to maximize the total expected payoff or return—over the entire n-day frost-protection season. On each night during the season, the orchardist must decide whether or not to take protective action. For simplicity, only two possible actions ($i = 0, 1$) are considered:

$$i = \begin{cases} 1 & \text{if protective action is taken,} \\ 0 & \text{if no protective action is taken.} \end{cases}$$

The cost of protection is taken to be a fixed amount, c say. If protective action is taken, then the temperature in the orchard is assumed to be raised by a constant amount Δ. Thus, if θ represents the air temperature in the environment surrounding the orchard and θ_e represents the "effective" air temperature in the orchard, then

$$\theta_e = \begin{cases} \theta + \Delta & \text{if} \quad i = 1, \\ \theta & \text{if} \quad i = 0. \end{cases}$$

The percent bud loss l is assumed to depend only on the effective temperature θ_e and to vary with the stage of bud development. Moreover, it is assumed that the relationship between percent bud loss and effective temperature can be characterized by a logistic function. Such a curve is more realistic than a straight line because the values of bud loss are constrained to range between 0 and 100%. Critical temperatures for 10 and 90% bud loss at each stage of development (Washington State University, 1971) are used to determine these logistic functions, which take the following form:

$$l = 100\%\{[\exp(\beta_0 + \beta_1\theta_e)] \times [1 + \exp(\beta_0 + \beta_1\theta_e)]^{-1}\}, \quad (1)$$

where β_0 and β_1 (with $\beta_1 < 0$) are parameters that vary with both the type of fruit and the stage of bud development. The percent bud loss l [specified by Eq. (1)] is multiplied by the percent buds remaining prior to the current night to obtain the net additional bud loss incurred on that night.

To convert bud loss into monetary expense, it is necessary to specify the relationship between percent bud loss l and percent yield loss y. This relationship is nonlinear, in part because it is possible to incur considerable bud loss and still obtain a full fruit crop. As a result, we take yield loss y to be 0% below a certain threshold value, l_t say, of percent bud loss and to increase from 0 to 100% as a quadratic function for values of bud loss between l_t and 100%. That is,

$$y = \begin{cases} 100\%[(l - l_t)/(100 - l_t)]^2 & \text{if} \quad l > l_t\,, \\ 0 & \text{if} \quad l < l_t\,. \end{cases} \quad (2)$$

Finally, a dollar value of fruit production per acre, d say, is used to translate percent yield loss into actual monetary loss.

c. *Meteorological information*

Three different types of meteorological information are considered: climatological information, imperfect forecasts (including currently available forecasts), and perfect forecasts. The observed minimum temperature on a given night can be viewed as a random variable Θ with probability density function $f_\Theta(\theta)$. The case of climatological information corresponds to estimating $f_\Theta(\theta)$ using historical temperature records. Similarly, the minimum temperature forecast for a given night can be viewed as a random variable Z with probability density function $f_Z(z)$. To characterize the accuracy of imperfect forecasts, the conditional distribution of the observed minimum temperature Θ, given a forecast $Z = z$, is also needed. This conditional distribution is assumed to be represented by a probability density function $f_{\Theta|Z}(\theta|z)$. If a Bayesian decision-analytic approach is taken, then $f_\Theta(\theta)$ and $f_{\Theta|Z}(\theta|z)$ represent the prior and posterior probability distributions of Θ. Here "prior" and "posterior" are interpreted as "before the forecast Z is received" and "after the forecast Z is received," respectively. Both climatological information and perfect forecasts can be considered to be limiting cases of imperfect forecasts. Specifically, climatological information corresponds to the situation in which $f_{\Theta|Z}(\theta|z)$ does not actually depend on z [i.e., $f_{\Theta|Z}(\theta|z) \equiv f_\Theta(\theta)$ for all z], whereas perfect forecasts correspond to the situation in which, given a forecast $Z = z$, the observed minimum temperature Θ equals z with probability one.

The three probability density functions $f_\Theta(\theta)$, $f_Z(z)$, and $f_{\Theta|Z}(\theta|z)$ are estimated using historical records of forecast and concomitant observed minimum temperatures. In this regard, only two of the three density functions need to be determined directly from the data because of the interrelationships among the densities. Since the chance of freezing temperatures decreases from the beginning to the end of the fruit-frost season, $f_\Theta(\theta)$ and $f_Z(z)$ are allowed to vary as a function of the calendar date. On the other hand, it is assumed that $f_{\Theta|Z}(\theta|z)$ does not vary with time. This assumption is supported by the forecast and observed minimum temperature data from Yakima, Washington (see Section 4a).

To estimate these probability density functions, it is convenient to first formulate a probabilistic model relating observed and forecast minimum temperatures. We assume that the joint distribution of Θ and Z is bivariate normal with parameters

$$\mu_\Theta = E(\Theta), \quad \mu_Z = E(Z), \quad \sigma_\Theta^2 = \mathrm{Var}(\Theta),$$
$$\sigma_Z^2 = \mathrm{Var}(Z), \quad \text{and} \quad \rho = \mathrm{Corr}(\Theta, Z),$$

where E denotes expected value, Var denotes variance, and Corr denotes correlation coefficient. Under the assumption of bivariate normality, $f_{\Theta|Z}(\theta|z)$ is a normal probability density function with parameters

$$E(\Theta|Z = z) = \mu_\Theta + \rho\frac{\sigma_\Theta}{\sigma_Z}(z - \mu_Z) \quad (3)$$

and

$$\mathrm{Var}(\Theta|Z = z) = \sigma_\Theta^2(1 - \rho^2) \quad (4)$$

(e.g., Hogg and Craig, 1970, p. 111).

The conditional expected value in Eq. (3) is a linear function of the forecast value z. Further, we assume that the slope of this line is equal to one; specifically, that

$$\rho\frac{\sigma_\Theta}{\sigma_Z} = 1. \quad (5)$$

If Eq. (5) holds, then Eqs. (3) and (4) become

$$E(\Theta|Z = z) = \mu_\Theta - \mu_Z + z \quad (6)$$

and

$$\mathrm{Var}(\Theta|Z = z) = \sigma_\Theta^2 - \sigma_Z^2\,, \quad (7)$$

respectively. With this simplified model, only the four parameters μ_Θ, μ_Z, σ_Θ^2, and σ_Z^2 need to be estimated.

A similar type of model was used by Phillips and Keelin (1978, Chapter 5) in their study of the value of long-range forecasts to the agricultural sector. The appropriateness of the assumptions required in the development of this model will be examined in Section 4a using historical records of observed and forecast minimum temperatures for Yakima, Washington.

One convenient property of this particular bivariate normal model is that the accuracy of the minimum temperature forecasts is completely characterized by the single conditional variance parameter given by Eq. (7). At least hypothetically, minimum temperature forecasts could be improved by reducing the value of this conditional variance. However, it is important to observe that, if the conditional density function $f_{\Theta|Z}(\theta|z)$ is adjusted, then the unconditional (or marginal) density function $f_Z(z)$ must also be adjusted to maintain consistency. Specifically, if we let Z^* denote an improved minimum temperature forecast and set

$$\mathrm{Var}(\Theta|Z^* = z) = \delta(\sigma_\Theta^2 - \sigma_Z^2), \quad (\delta > 0) \tag{8}$$

then it can be easily shown, using Eqs. (6), (7) and (8), that $f_{Z^*}(z)$ is a normal probability density function with parameters

$$E(Z^*) = \mu_Z \tag{9}$$

and

$$\mathrm{Var}(Z^*) = \delta\sigma_Z^2 + (1 - \delta)\sigma_\Theta^2 . \tag{10}$$

It should be noted that a δ value of one corresponds to the case of current imperfect forecasts, a δ value of zero to perfect forecasts, and values of δ between zero and one to improved forecasts. Since Eqs. (9) and (10) hold even if δ is greater than one, these expressions also can be used to investigate forecasts that are actually of poorer quality than current forecasts. As a practical upper bound on δ, climatological forecasts can be shown to be equivalent to the case in which

$$\delta = \sigma_\Theta^2/(\sigma_\Theta^2 - \sigma_Z^2). \tag{11}$$

d. Markov decision processes and the value of information

As indicated in Section 3b, the goal of the orchardist is to minimize, on the basis of the available meteorological information, the total expected expense over the entire n-day frost-protection season. Because current actions and outcomes are related to both previous and future actions and outcomes, this decision-making problem is dynamic in nature. For the purposes of making current decisions, we assume that previous actions and outcomes can be completely characterized in terms of a single random variable. This variable is the total percent bud loss, denoted by L, accumulated through the season prior to the current day. On a given day in the season, the orchardist uses both the minimum temperature forecast $Z = z$ and the total previous bud loss state $L = l$ in deciding whether or not to take protective action on that night. The optimal action, given $Z = z$ and $L = l$, is that action i that leads to minimizing the expected expense over the remainder of the season. Because of the dynamic nature of the decision-making problem and because of the nonstationarity of both the climatological probability distribution of minimum temperature and of the relationship between bud loss and minimum temperature, the optimal actions also vary with the day of the season.

Under these assumptions, a particular class of sequential decision-making models, known as Markov decision processes (e.g., Howard, 1960; Ross, 1970), can be applied to the fruit-frost situation. In the case of a Markov decision process, a computational technique called dynamic programming (e.g., White, 1978) can be used to identify the optimal actions and to calculate the associated minimum expected expenses. First the optimal actions and minimum expected expenses are determined, as a function of the forecast z and previous bud loss l, on the last day of the frost-protection season. Then, the results for the last day of the season are employed, by means of a recurrence relation, to determine the optimal actions and minimum expected expenses on the next to last day of the season. The iterative procedure, sometimes called "backwards induction," is continued until the first day of the season is reached. As a result, the expected expenses can be computed for the entire n-day frost-protection season. In this procedure, the expected expenses associated with climatological information, imperfect forecasts, and perfect forecasts are denoted by EC, EF, and EP, respectively. The specific recurrence relations employed in determining these expected expenses are given in the Appendix.

The value of meteorological information is considered to be the reduction in expected expense from the situation in which such information is not available to a situation in which it is available. Since it is reasonable to assume that climatological information would always be available to the orchardist, the value of the different types of meteorological information is measured relative to that of climatological information. On this scale, climatological information has zero value by definition, whereas the value of imperfect forecasts (VF) and of perfect forecasts (VP) are given by

$$VF = EC - EF \tag{12}$$

and

$$VP = EC - EP, \tag{13}$$

in which EC, EF, and EP are calculated using Eqs. (A1), (A2), and (A3), respectively (see Appendix). Because an *ex ante* approach to assessing the value of information has been taken,

$$0 \leqslant VF \leqslant VP. \quad (14)$$

In other words, the value of imperfect forecasts must be both greater than or equal to the value of climatological information and less than or equal to the value of perfect forecasts. A proof of Eq. (14) for Markov decision processes is presented in Katz *et al.* (1981), and this relationship can be shown to hold, in general, for any decision process.

4. Results

In this section the results of a study of the economic value of frost forecasts to orchardists in the Yakima Valley are presented. The case of red delicious apples is considered in detail, with forecast value estimates provided for this variety of fruit as well as for bartlett pears and elberta peaches. These value estimates are based on the application of the dynamic decision-making model discussed in Section 3. First the values of the parameters of the Markov decision process are specified. Because of the dynamic nature of the decision-making problem, the optimal policy (i.e., the rule for choosing actions) for the orchardist varies throughout the frost-protection season, and examples of the optimal policy on selected calendar dates during the frost season are presented. Then the estimates of the value of different types of meteorological information are given, including changes in the value estimates as a function of the accuracy of imperfect forecasts.

a. Parameter estimates

The values of the parameters of the Markov decision process will be specified for the case of red delicious apples. Parameter values expressed in monetary terms correspond to the year 1977. The frost-protection season is taken as the time period from 16 March–30 May, so that $n = 76$. We have assumed a protection effect of $\Delta = 5°F$, as in Baquet *et al.* (1976), and a cost of protection based on the use of heaters of $c = \$70$ per acre-day (J. K. Ballard, Yakima County Extension Agent, pers. comm., 1977). The estimates of the parameters β_0 and β_1 of the logistic functions [see Eq. (1)] representing the relationship between percent bud loss and minimum temperature are listed in Table 1. These estimates were obtained for specific time periods, based on average dates, for the different stages of bud development. Fig. 1 shows the logistic functions corresponding to these parameter values. The threshold value l_t of percent bud loss below which no yield loss occurs, needed in the relationship between yield loss and bud loss [see Eq. (2)], was set equal to 50%. For the monetary value of fruit production d, the 1977 dollar value of \$2348 for production per acre of apples in the Yakima Valley was used. This estimate includes red delicious as well as other varieties of apples and is based on data from the Yakima Project (Bureau of Reclamation, 1977), an irrigation project covering the Yakima Valley.

TABLE 1. Parameters of logistic functions [see Eq. (1)] relating bud loss (%) to minimum temperature (°F) for red delicious apples.

Time period	β_0	β_1
16–23 March	7.69	−0.549
24–31 March	10.44	−0.549
1–5 April	17.58	−0.732
6–9 April	28.56	−1.099
10 April–30 May	38.82	−1.465

Observed and forecast minimum temperatures at the Yakima key station during the frost-protection season (at most 16 March–30 May) for the 20-year period 1957–76 were obtained from published records (National Weather Service, 1957–76). Fig. 2 shows some descriptive statistics for the conditional distribution of observed minimum temperature Θ, given a forecast $Z = z$. The descriptive statistics shown were derived by first determining the conditional lower quartile, median, and upper quartile for each possible value of z and then smoothing these conditional statistics using a technique called hanning (Tukey, 1977, p. 231).

The data presented in Fig. 2 have several characteristics that are consistent with the assumption of a bivariate normal distribution for Θ and Z made in Section 3c. As indicated by the conditional medians, the conditional mean appears to be a linear function of z in agreement with Eq. (3). The con-

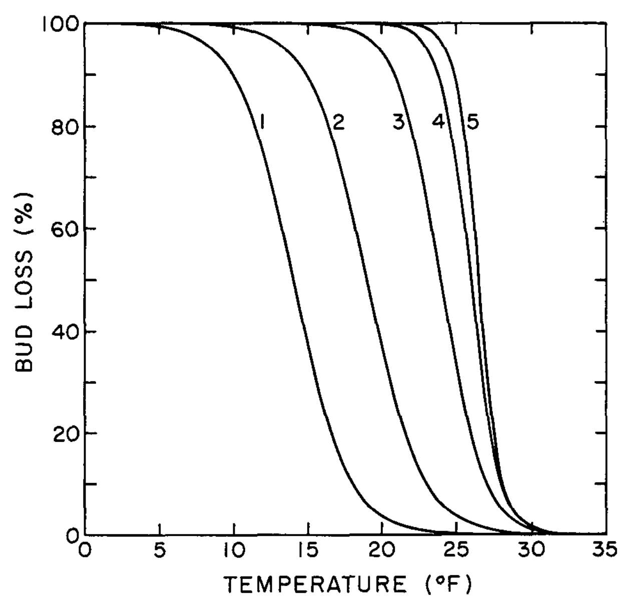

FIG. 1. Logistic functions relating bud loss and minimum temperature for red delicious apples for the time periods: 1 = 16–23 March, 2 = 24–31 March, 3 = 1–5 April, 4 = 6–9 April, 5 = 10 April–30 May.

ditional interquartile ranges (i.e., the differences between the upper and lower quartiles) suggest that the conditional variance is constant (rather than dependent upon the value of z) and thereby consistent with Eq. (4). Since the slope of the conditional medians appears to be approximately equal to one, the simplification represented by Eq. (5) is also reasonable.

The assumption that the conditional density function $f_{\Theta|Z}(\theta|z)$ does not vary throughout the fruit-frost season was checked by comparing conditional distributions for the first half and second half of the season, using the Yakima minimum temperature data. The correspondence between these two conditional distributions appears to be quite good. An alternative approach would be to assume that the conditional distribution of the forecast minimum temperature Z given the observed minimum temperature $\Theta = \theta$, with associated density function $f_{Z|\Theta}(z|\theta)$, remains constant throughout the season. We note that if $f_{\Theta|Z}(\theta|z)$ remains constant then $f_{Z|\Theta}(z|\theta)$ cannot remain constant and vice versa. The Yakima minimum temperature data indicate that the difference between the conditional distributions of Z given Θ for the first half and second half of the season was much greater than that between the two conditional distributions of Θ given Z.

The model presented in Section 3 requires the estimation of the parameters μ_Θ, μ_Z, σ_Θ^2, and σ_Z^2 associated with the density functions $f_\Theta(\theta)$ and $f_Z(z)$. These parameters are allowed to change each half month of the frost-protection season (i.e., every 15–16 days). Table 2 lists the parameter values for the

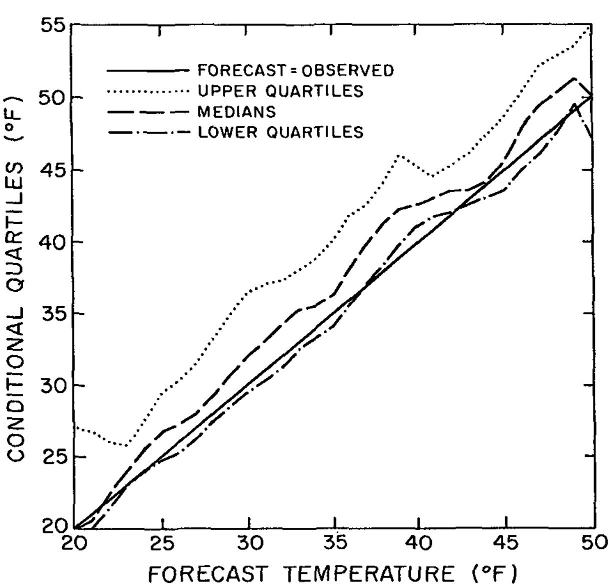

FIG. 2. Smoothed lower quartiles, medians, and upper quartiles for the conditional distribution of observed minimum temperature given forecast minimum temperature at Yakima, Washington.

TABLE 2. Parameters of probability distributions of observed and forecast minimum temperatures at Yakima, Washington.

Time period	Observed temperatures (°F) μ_Θ	σ_Θ	Forecast temperatures (°F) μ_Z	σ_Z
16–31 March	32	6.5	31	5.6
1–15 April	34	6.4	32	5.1
16–30 April	36	6.5	33	5.4
1–15 May	41	6.9	38	6.1
16–30 May	43	7.1	41	6.3

various time periods, based on sample means and standard deviations obtained from the 20-year Yakima record. The parameter estimates associated with the density function $f_{\Theta|Z}(\theta|z)$ [see Eqs. (6) and (7)] are of the form

$$E(\Theta|Z = z) = 2.5°\text{F} + z \qquad (15)$$

and

$$\text{Var}(\Theta|Z = z) = (3.5°\text{F})^2. \qquad (16)$$

Because this density function is not allowed to vary throughout the frost-protection season, the parameter estimates were derived by averaging the appropriate means and variances for the half-month time periods. In particular, the conditional variance parameter in Eq. (7) that completely characterizes the accuracy of currently available Yakima minimum temperature forecasts is equal to $(3.5°\text{F})^2$ [Eq. (16)]. We note that, as is allowed for in Eq. (6) and as is evident in Fig. 2, a consistent tendency to underforecast the minimum temperature by ~2.5°F on the average is present [Eq. (15)]. Since the Yakima key station is located in an area that tends to have lower minimum temperatures than many of the orchards in the Yakima Valley, the estimates of forecast value presented here were based on climatological means taken to be 2°F higher than those listed in Table 2.

b. Value estimates

RED DELICIOUS APPLES. Using the parameter estimates corresponding to the case study of red delicious apples in the Yakima Valley, the dynamic decision-making model discussed in Section 3 was employed to assess the economic value of current minimum temperature forecasts and perfect forecasts to orchardists. Before discussing the value estimates that were obtained, we provide some examples of the nature of the orchardist's optimal policy to demonstrate the dynamic nature of the fruit-frost decision-making situation.

For the first and last days of the frost-protection season, together with another day in the middle of the season, Fig. 3 shows the values of the minimum

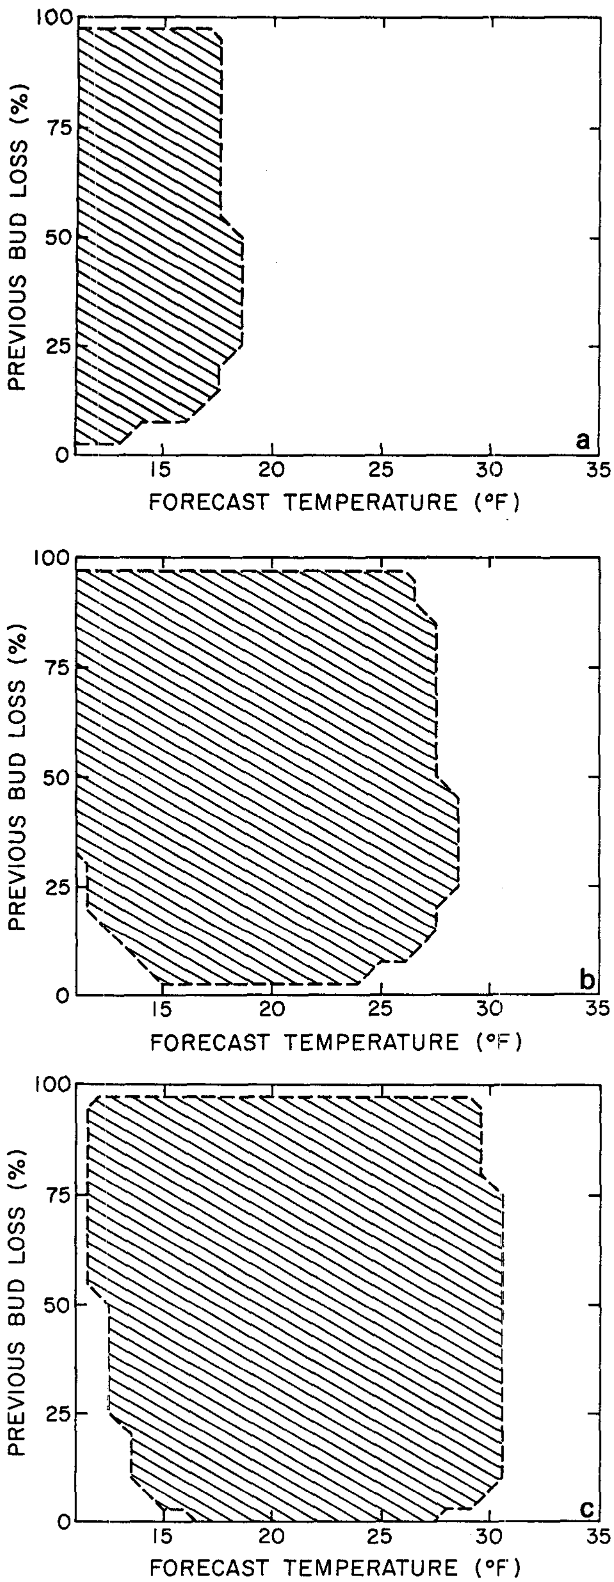

FIG. 3. Optimal policy for the orchardist as a function of previous total bud loss and forecast minimum temperature on (a) day 1, (b) day 21, and (c) day 76. The optimal action is to protect inside shaded region and not to protect outside shaded region.

temperature forecast $Z = z$ and of the previous percent bud loss $L = l$ for which protection ($i = 1$) is the optimal action. These regions were determined for each of the 76 days of the frost-protection season in the course of the dynamic programming algorithms using Eqs. (A1), (A2), and (A3) (see Appendix). For example, on day 21 of the season (Figure 3b), if 50% bud loss has previously occurred (i.e., $L = 50\%$), then the orchardist should take protective action whenever the forecast minimum temperature is less than or equal to 27°F (i.e., $Z \leqslant 27°F$). On the other hand, if 50% bud loss has occurred prior to the last day of the season, then the orchardist should take protective action on the last day (Figure 3c) whenever the forecast minimum temperature is between 12 and 30°F (i.e., $12°F \leqslant Z \leqslant 30°F$). Judging from Fig. 3, it is evident that protection is optimal for a much wider range of values of forecast minimum temperature and percent bud loss at the end of the season than at the beginning of the season. When the nature of the optimal policy is examined for all 76 days of the frost-protection season (detailed results not included in this paper), it is also evident that the shape of the region for which protection is optimal changes from day to day quite rapidly near the beginning of the season and only slowly near the end of the season.

The expected expenses for climatological information, current imperfect forecasts, and perfect forecasts were calculated by employing the dynamic programming algorithms specified in Eqs. (A1), (A2), and (A3). These expected expense estimates are included in Table 3, together with the value estimates for current and perfect forecasts obtained from the expected expenses by applying Eqs. (12) and (13). The value estimates are expressed in terms of dollars per acre over the entire frost-protection season for the year 1977. The results in Table 3 indicate that the value of current forecasts is estimated to be \$808 per acre, whereas the value of perfect forecasts is about \$1233 per acre. Comparing these two value estimates, we see that current forecasts have already realized ~66% of the potential reduction in expected expense over climatological information.

OTHER TYPES OF FRUIT. In addition to the case study of red delicious apples, estimates of the economic value of minimum temperature forecasts to orchardists in the Yakima Valley were also obtained

TABLE 3. Expected expense and value estimates for red delicious apples. Monetary value of production d equals \$2348 per acre.

Type of information	Expected expense (\$/acre)	Value (\$/acre)
Climatological information	1721	0
Current forecasts	913	808
Perfect forecasts	488	1233

for both bartlett pears and elberta peaches. The same dynamic decision-making methodology was applied, with some changes in the parameter values of the Markov decision process. Because the definition of the stages of bud development, the average dates on which they occur, and the critical minimum temperatures corresponding to 10 and 90% bud loss all vary with either the type or variety of fruit (Washington State University, 1971), the parameters of the logistic functions that relate percent bud loss to minimum temperature [see Eq. (1)] are different for both bartlett pears and elberta peaches (the specific logistic parameter values for these two cases are not included in this paper). The monetary value of production also depends on the type of fruit and was taken to be d = \$1856 per acre for pears and d = \$1240 per acre for peaches (Bureau of Reclamation, 1977).

Tables 4 and 5 include estimates of the economic value of both current and perfect minimum temperature forecasts for bartlett pears and elberta peaches, respectively. These estimates are also expressed in terms of dollars per acre over the entire frost-protection season for the year 1977. As indicated by Table 4, the value of current forecasts is estimated to be \$492 per acre for bartlett pears, or ~63% of the value of perfect forecasts. Similarly, Table 5 reveals that the value of current forecasts is about \$270 per acre for elberta peaches, or ~47% of the value of perfect forecasts. Comparing these two percentages with that for red delicious apples (i.e., 66%), we see that in the cases of red delicious apples and bartlett pears about the same percentage of potential value has been realized, although the value estimates for these two types of fruit are quite different. Elberta peaches, on the other hand, have substantially smaller value estimates, both in absolute and percentage terms.

c. *Comparisons of accuracy and value*

The accuracy of minimum temperature forecasts, assuming the bivariate normal model introduced in Section 3c, is characterized by a single conditional variance parameter [see Eq. (7)]. Equivalently, the accuracy of forecasts can be summarized in terms of the conditional standard deviation, or square root of the conditional variance in Eq. (7), to provide a measure of accuracy in the same units as the meteorological variable being forecast (i.e., in °F). The conditional standard deviation of current minimum temperature forecasts in Yakima is 3.5°F, whereas the conditional standard deviation for climatological information is ~6.7°F (derived by averaging the climatological variances for the five time periods given in Table 2). Since the conditional standard deviation of perfect forecasts is by definition equal to zero, current minimum temperature forecasts have attained ~52% of the potential improvement in accuracy over climatological information. Comparing this percentage improvement in forecast accuracy to the percentage increases in the economic value of forecasts given earlier for each of the three types of fruit, we observe that percentage increases in forecast value for both red delicious apples and bartlett pears are greater than the percentage improvement in forecast accuracy. The percentage increase in forecast value for elberta peaches, in contrast, is less than the percentage improvement in forecast accuracy.

TABLE 4. Expected expenses and value estimates for bartlett pears. Monetary value of production d equals \$1856 per acre.

Type of information	Expected expense (\$/acre)	Value (\$/acre)
Climatological information	1043	0
Current forecasts	551	492
Perfect forecasts	256	787

TABLE 5. Expected expenses and value estimates for elberta peaches. Monetary value of production d equals \$1240 per acre.

Type of information	Expected expense (\$/acre)	Value (\$/acre)
Climatological information	868	0
Current forecasts	597	270
Perfect forecasts	299	569

These results suggest that a nonlinear relationship exists between the accuracy and value of meteorological information. To investigate the nature of this relationship in greater detail, the methodology discussed in Section 3c for hypothetically improving current forecasts was employed. The parameter δ, defined in Eq. (8), was set equal to $(1/3)^2$, $(2/3)^2$, and $(4/3)^2$, corresponding to situations in which the conditional standard deviation of the new forecasts is one-third, two-thirds, and four-thirds, respectively, of that for current forecasts. We note that the value of $\delta = (4/3)^2$ corresponds to a case of forecasts of poorer quality than current forecasts, but of better quality than climatological information because the constraint imposed by the condition in Eq. (11) is satisfied.

For each type of fruit, the economic value of forecasts corresponding to each of the three δ values was obtained. Using these value estimates, together with those estimates already obtained for climatological information, current forecasts, and perfect forecasts, curves of forecast value versus forecast accuracy were constructed. To make these curves directly comparable, the estimates of forecast value for each type of fruit were divided by the associated monetary value of production d, and the curves are shown in this relative form in Fig. 4. All three curves exhibit approximately the same degree of nonlinearity. Because the absolute value of the slope of each curve can be interpreted as the incremental increase in

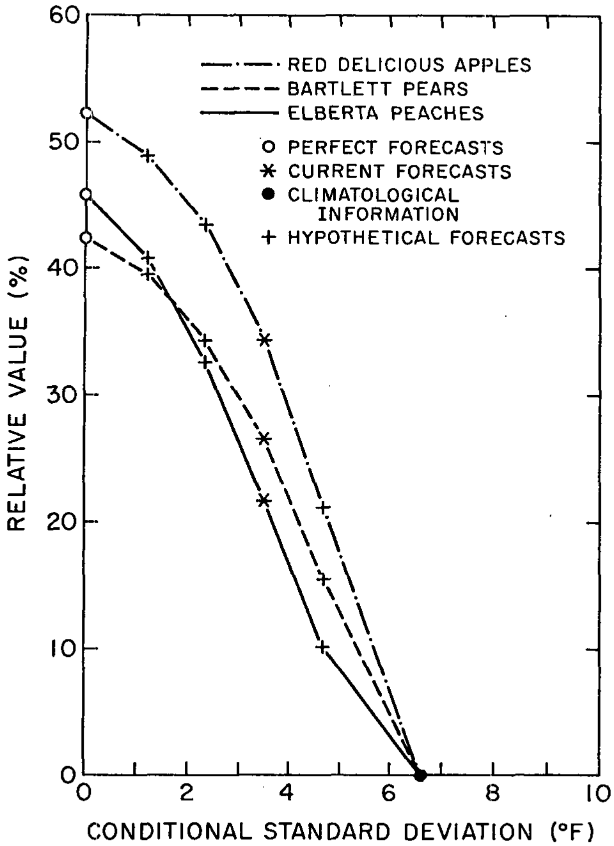

FIG. 4. Economic value of minimum temperature forecasts to the orchardist (divided by total value of fruit production) as a function of the conditional standard deviation of the observed minimum temperature given the forecast minimum temperature for red delicious apples, bartlett pears, and elberta peaches.

forecast value associated with a unit improvement in forecast accuracy, the nature of the nonlinearity is such that continued improvements in accuracy over current forecasts correspond to gradually smaller increases in forecast value. Hence, improvements in current forecasts will not be as significant in terms of economic value to orchardists as were comparable improvements in the past.

5. Summary and conclusion

In this paper we have described the results of a study in which a dynamic decision-making model was used to assess the value of frost (i.e., minimum temperature) forecasts to orchardists in the Yakima Valley of central Washington. The problem of determining the value of meteorological information in this context was approached from a decision-analytic point of view. A dynamic model was formulated to describe the relationships between decisions and events on different occasions and to provide estimates of the value of information that reflect the orchardist's overall objective of minimizing expected expenses (or maximizing expected payoffs) over the entire frost-protection season. This model constitutes a Markov decision process, and a computational technique known as dynamic programming was used to identify the orchardist's optimal actions over the season and to determine the associated minimum expected expenses.

Some results concerning the value of forecasts to orchardists were presented for the cases of red delicious apples, bartlett pears, and elberta peaches, with the expected expense corresponding to climatological information representing the zero point on the value scale. Over the entire frost-protection season, these value estimates (in 1977 dollars) were \$808 per acre for red delicious apples, \$492 per acre for bartlett pears, and \$270 per acre for elberta peaches. These amounts account for 66, 63 and 47%, respectively, of the economic value that would be realized by an orchardist whose decisions were based on perfect information. An investigation of the relationship between the accuracy and value of these minimum temperature forecasts revealed that this relationship is nonlinear and that improvements in current forecasts would not be as significant to orchardists in an economic sense as were comparable improvements in the past.

Decision analysis provides a particularly appropriate framework within which to study the value of meteorological information. The components of this framework were briefly described in Section 3a. Despite the obvious strengths of this methodology, it is important to mention two issues that require further attention when decision analysis methodology is used in practice. First, as indicated in Section 3a, decision analysis is prescriptive rather than descriptive in nature. That is, the objective of decision analysis is to prescribe how decisions can be made and information can be used in a rational, optimal manner, not to describe how decisions actually are made and how information is used in practice. Since the application of decision analysis inevitably requires the analyst to make certain assumptions regarding (*inter alia*) the preferences and/or behavior of the decision maker, it is important to compare the results of a formal prescriptive analysis of a decision-making problem with the decision maker's actual practices and procedures. In the fruit-frost problem, for example, it would be desirable to investigate the role that information from frost alarms (automatic devices that alert the fruit grower to the occurrence of critical temperatures in the orchard) and from other sources plays in the orchardist's decision-making process and the impact of such information on the estimates of the value of the specialized NWS temperature forecasts.

Second, decision analysis focuses on an individual decision maker and provides a means of determining the value of information to such an individual. However, in many meteorological situations such as the fruit-frost problem, the information provided by NWS forecasters is used by many individuals to make identical or similar decisions. Obviously, it would be desirable to be able to aggregate the individual value estimates in some manner over the entire user group to obtain an overall estimate of the value of the information to all orchardists. The existence of possible market and price effects suggests that a realistic overall estimate cannot be obtained simply by multiplying the value of the information to an individual orchardist by the number of orchardists. Unfortunately, decision analysis itself provides no straightforward solution to this problem. Thus, the problem of aggregating value estimates over user groups will require further attention in the future.

We believe that the use of the dynamic decision-making model developed in this paper to assess the value of meteorological information in the fruit-frost situation represents a significant advance over previous studies of this problem. As previously indicated, the static model employed by Katz and Murphy (1979) could be used only to estimate the value of minimum temperature forecasts to orchardists on a specific calendar date during the frost season. In this case, it was assumed that the orchardist wanted to minimize the expected expense on that occasion and any previous bud loss that may have been incurred was ignored. In the study by Baquet *et al.* (1976), the orchardist was assumed to want to maximize the expected payoff or return on a particular occasion taking into account the previous bud loss, and their solution to the value-of-information problem involved the use of a simulation procedure.

In the model developed in this paper, the orchardist is assumed to want to minimize expected expense (or, equivalently, maximize expected payoff) over the entire frost-protection season, which leads to the formulation of a fully dynamic model. An exact numerical solution to this decision-making problem is obtained by dynamic programming. Moreover, the model of meteorological information employed in this paper enables us to investigate the value of such information as a function of its quality (including cases of improved information). We hasten to add, however, that many possible extensions of the work described here on the fruit-frost problem are possible. For example, we could consider a larger and more realistic set of possible actions, including the use of wind machines and overhead sprinklers in conjunction with orchard heaters as well as actions involving partial protection. Another important extension would be to assess and/or model the utility functions of individual orchardists so that their attitudes toward risk could be properly considered in modeling the decision-making problem and in estimating the value of meteorological information. The impact of utility functions in this context was studied by Baquet *et al.* (1976). Another possible extension would be to consider different types of meteorological information, including minimum temperature forecasts for the second day, the third day, etc., as well as for the first day. In addition, it would be of interest to investigate the situation in which the short-term (tactical) decision-making problem studied in this paper was embedded in the long-term (strategic) problem of which types of protective devices to purchase and what kinds of other resources to have available. Such a study would involve the consideration of both short-term forecasts and climatological data.

Finally, it would be of considerable interest to investigate further extensions of the methodology developed in this paper as well as the application of this methodology to other weather and/or climate information-sensitive problems. With regard to the former, the relationships among the expected values of different types of information in Markov decision processes are described by Katz *et al.* (1981) and the application of this methodology to an extension of the familiar cost-loss ratio situation will be described in a forthcoming paper. The dynamic model developed in this paper could be applied to a variety of decision-making problems of interest to meteorologists, such as irrigation scheduling and ship routing, and we plan to study these and other possible applications in the future.

Acknowledgments. We would like to thank Mr. J. K. Ballard (County Extension Agent, Yakima, WA) and Dr. E. L. Proebsting, Jr. (Washington State University) for providing us with invaluable data and background information on the fruit-frost problem in the Yakima Valley. We would also like to express our appreciation to Mr. C. B. Graves, Jr. (National Weather Service) for making available the annual reports of the NWS fruit-frost forecasting project. Special thanks are due to Mr. W.-R. Hsu who provided valuable programming assistance during the latter stages of the work reported in this paper.

APPENDIX

Dynamic Programming Algorithms

Recurrence relations are required to compute the total expected expense for the entire n-day frost-protection season for climatological information (EC), for imperfect forecasts (EF), and for perfect forecasts (EP). We restrict the random variables Θ and

Z to integer values and let $P_\Theta(\theta) = \Pr\{\Theta = \theta\}$, the climatological probability distribution of observed minimum temperature, $P_Z(z) = \Pr\{Z = z\}$, the probability distribution of minimum temperature forecasts, and $P_{\Theta|Z}(\theta|z) = \Pr\{\Theta = \theta|Z = z\}$, the conditional distribution of observed minimum temperature given the minimum temperature forecast z. These discrete distributions are obtained by applying the so-called continuity correction to the corresponding normal probability density functions; for example,

$$P_\Theta(\theta) = \int_{\theta-1/2}^{\theta+1/2} f_\Theta(\theta')d\theta'.$$

The immediate expense incurred by taking action i is denoted by $CP(i)$. In particular, $CP(0) = 0$ and $CP(1) = c$. The terminal (or end of season) expense corresponding to total percent bud loss l is denoted by $CL(l)$. This terminal expense is obtained by applying Eq. (2), determining the percent yield loss y corresponding to percent bud loss l, and then multiplying y by the monetary value of production d. The transfer function $T(\theta, l, i)$ provides the total percent bud loss through a given day when the observed minimum temperature on that day is θ, the total percent bud loss prior to that day is l, and action i is taken on that day. This transfer function is obtained by applying the logistic functions that relate bud loss to minimum temperature [see Eq. (1)]. For notational simplicity, the dependency of $P_\Theta(\theta)$, $P_Z(z)$, $P_{\Theta|Z}(\theta|z)$, and $T(\theta, l, i)$ on the day of the season has been omitted.

Climatological information

We let $EC_k(l)$ denote the total expected expense for the last k days of the season based on taking optimal actions (i.e., minimizing the total expected expense for the remainder of the season) if climatological information is available and conditional on total percent bud loss l prior to the last k days. Since no bud loss should have occurred prior to the first day of the season, the desired quantity EC is simply given by $EC_n(0)$. The dynamic programming algorithm is specified by the recurrence relation

$$EC_k(l) = \min_{i=0,1} \{CP(i) + \sum_\theta P_\Theta(\theta)EC_{k-1}[T(\theta, l, i)]\}, \quad \text{(A1)}$$

$k = 1, 2, \ldots, n$, with the convention that $EC_0(l) = CL(l)$.

Imperfect forecasts

We let $EF_k(l)$ denote the total expected expense for the last k days of the season based on taking optimal actions if imperfect forecasts are available and conditional on total percent bud loss l prior to the last k days. In this notation, $EF = EF_n(0)$. The dynamic programming algorithm is specified by the recurrence relation

$$EF_k(l) = \sum_z P_Z(z) \min_{i=0,1} \{CP(i) + \sum_\theta P_{\Theta|Z}(\theta|z)EF_{k-1}[T(\theta, l, i)]\}, \quad \text{(A2)}$$

$k = 1, 2, \ldots, n$, with the convention that $EF_0(l) = CL(l)$.

Perfect forecasts

We let $EP_k(l)$ denote the total expected expense for the last k days of the season based on taking optimal actions if perfect forecasts are available and conditional on total percent bud loss l prior to the last k days. In this notation, $EP = EP_n(0)$. The dynamic programming algorithm is specified by the recurrence relation

$$EP_k(l) = \sum_\theta P_\Theta(\theta) \min_{i=0,1} \{CP(i) + EP_{k-1}[T(\theta, l, i)]\}, \quad \text{(A3)}$$

$k = 1, 2, \ldots, n$, with the convention that $EP_0(l) = CL(l)$.

Details on computations

Although the percent bud loss state l is a continuous variable, to make the computations feasible only a finite number of possible values, say $\{l_1, l_2, \ldots, l_m\}$, are permitted in practice. For each value of k, the recurrence relations must be evaluated for all possible values of l; that is, $l = l_1, l_2, \ldots, l_m$. The computations proceed in an iterative manner through repeated application of the appropriate recurrence relation, either Eq. (A1), Eq. (A2), or Eq. (A3), by setting $k = 1, 2, \ldots, n$. The actual calculations were performed with l restricted to the set of 21 values $\{0, 5, \ldots, 100\%\}$.

REFERENCES

Anderson, J. R., J. L. Dillon and B. Hardaker, 1977: *Agricultural Decision Analysis.* Iowa State University, Ames, 344 pp.

Bagdonas, A., J. C. Georg and J. F. Gerber, 1978: Techniques of frost prediction and methods of frost and cold protection. WMO, Tech. Note No. 157, 160 pp.

Baquet, A. E., A. N. Halter and F. S. Conklin, 1976: The value of frost forecasting: a Bayesian appraisal. *Amer. J. Agric. Econ.,* **58,** 511–520.

Bureau of Reclamation, 1977: Yakima Project crop production report. Yakima, WA (unpub. rep.), 2 pp.

Hogg, R. V., and A. T. Craig, 1970: *Introduction to Mathematical Statistics,* 3rd ed. Macmillan, 415 pp.

Howard, R. A., 1960: *Dynamic Programming and Markov Processes.* MIT Press, 136 pp.

Katz, R. W., and A. H. Murphy, 1979: Assessing the value of frost forecasts to orchardists: a decision analytic approach. *Preprints Fourteenth Conf. Agriculture and Forest Meteorology*, Minneapolis, Amer. Meteor. Soc., 34–37.

Katz, R. W., A. H. Murphy and R. L. Winkler, 1981: Markov decision processes and the value of meteorological information. *Preprints Seventh Conf. Probability and Statistics in Atmospheric Sciences*, Monterey, Amer. Meteor. Soc., 125–128.

Keeney, R. L., 1982: Decision analysis: state of the field. *Oper. Res.*, **30**, (in press).

Lindley, D. V., 1971: *Making Decisions*. Wiley-Interscience, 309 pp.

National Weather Service, 1957–1976: Report of fruit-frost project, Yakima and Benton Counties, Washington. U.S. Dept. Commerce, NOAA, National Weather Service, annual reports, Yakima, WA.

Phillips, L. D., and T. W. Keelin, 1978: Bayesian modelling of improved climatological forecasts for large agricultural models. *Multidisciplinary Research Related to the Atmospheric Sciences*, M. H. Glantz, H. van Loon, and E. Armstrong, Eds., National Center for Atmospheric Research, 213–233.

Raiffa, H., 1968: *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Addison-Wesley, 309 pp.

Ross, S. M., 1970: *Applied Probability Models with Optimization Applications*. Holden-Day, 198 pp.

Tukey, J. W., 1977: *Exploratory Data Analysis*. Addison-Wesley, 688 pp.

Washington State University, 1971: Apples—Critical temperatures for blossom buds. Pullman, Washington State University, Research and Extension Service, pamphlet, 2 pp. Similar pamphlets are available for pears and peaches (also 2 pp. each).

White, D. J., 1978: *Finite Dynamic Programming*. Wiley-Interscience, 204 pp.